LA-UR -82-83

CONF-820902--1

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

LA-UR--82-83

DE82 010412

MASTER

TITLE: BRINGING A LARGE COMPUTER NETWORK INTO FOCUS

AUTHOR(S): Norman R. Morse, C-DO
Joseph L. Thompson, C-8

SUBMITTED TO: 6th International Conference on Software Engineering, September 13-16, 1982, Tokyo, Japan.

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

BRINGING A LARGE COMPUTER NETWORK INTO FOCUS

by

Norman R. Morse and Joseph L. Thompson
Los Alamos National Laboratory
Los Alamos, New Mexico, USA

ABSTRACT

In this paper we describe the development and implementation of the Facility for Operations Control and Utilization Statistics (FOCUS), a new centralized node in the Integrated Computing Network of the Los Alamos National Laboratory. FOCUS consists of production control, performance measurement, and network information subsystems. The software engineering practices on which the development was based are discussed, with emphasis on the application of those practices to network systems development.

Keywords: structured development; distributed operating systems; production control; network performance; network status and information; centralized operator stations; network systems development.

I. INTRODUCTION

A computing facility can no longer be expected to consist of a general purpose mainframe system supplying all of the computational requirements of an organization. Due to the decreasing costs and increasing power of computing equipment, computing facilities are evolving toward network systems consisting of user terminals, network communications, mainframe processors, and specialized satellite functions such as file storage and output media stations.

Software engineering for such network systems present both opportunities (for example, improved modularity) and challenges (for example, complex communications). This paper describes the development and implementation of a new production control node for the Integrated Computing Network (ICN) of the Los Alamos National Laboratory. The new node is called the Facility for Operations Control and Utilization Statistics (FOCUS), and it serves as a central network platform for controlling production tasks on mainframe processors, collecting network status and performance data, and providing network information to both operators and network users. The development of FOCUS was based on concepts and software engineering practices currently in use by the Laboratory's Computer Systems Group.

II. BACKGROUND

Located at an altitude of 7500 feet in the mountains of Northern New Mexico, the Los Alamos National Laboratory is a research and development facility with emphasis on nuclear research and energy programs. The Laboratory is managed by the University of California for the United States Department of Energy. Of its approximately 7,000 employees, about 2,200 are technical staff members in physics, engineering, mathematics, and computer science.

Programmatic requirements have dictated for many years that the Laboratory maintain an immense computational capacity. Mainframe processors include three CDC 6600s, four CDC 7600s, and four Cray-1 supercomputers. Figure 1 illustrates the current network. It evolved from a batch facility in the early '70s into a large network currently serving more than 3,000 interactive users. During the past few years, building on our network experience, we separated two major operating system functions and developed them as specialized internal nodes in the network. They were the Print and Graphics Express Station (PAGES)[1] and the Common File System (CFS).[2]

The PAGES system handles virtually all printed and graphical output from the ICN. PAGES features two XEROX 9700 laser printers, three FR-80 microfilm recorders, and several other output devices controlled by dual VAX-11/780 controllers.

The CFS is a centralized file system for the network based on IBM 3850 Main Storage System technology managed by dual IBM 4341 controllers. Users gain access to the CFS node through utility software on the mainframe processors, allowing them to store, retrieve and manage their files. CFS storage is hierarchical, with frequently accessed files on IBM 3350 and 3380 disks, less frequently accessed files on cartridges in the IBM 3850 Mass Storage System, and archival or low-use files kept on 3850 cartridges stored in offline cabinets. The CFS features an innovative file migration program that automatically moves files between disk and online/offline mass storage based on each file's usage and size. Development of the CFS was based on the structured techniques of Yourdon[3] and DeMarco.[4] CFS was delivered ahead of schedule and has proven extremely reliable; the software has not lost or damaged one of its 450,000 files since installed in 1978.

The Los Alamos ICN has continued to migrate toward distributed computing. Two dozen VAX-11/780 minicomputers serve as distributed processors throughout the Laboratory, providing several hundred users a local capability backed by CFS, PAGES, and the 11 mainframe processors mentioned previously. In addition, intelligent terminals to be used as scientific work stations are currently under development.

As a more distributed network evolves, users need a more general way of submitting production jobs to the mainframe processors and of determining the status of each network node. With three major operating systems running on 11 mainframes, automated load leveling across processors is required to balance work and to handle periods of processor unavailability.

III. FOCUS OBJECTIVE

The objective of FOCUS is to enhance the effectiveness and reliability of the Los Alamos ICN. FOCUS is an ICN node that provides a centralized point for the control of mainframe processor jobs, for the display of network status and utilization information, and for the collection of performance monitoring data from network nodes (including mainframes).

The ICN grows more complex as new classes of computers are integrated into the network. To offset this complexity we are introducing conceptual simplifications by concentrating duplicated functions into centralized, internal nodes of the network. Before FOCUS, both control and status information were very decentralized. Previously, each mainframe ran its own production or "batch" subsystem, and the production user interface was different for each of the three operating systems in use. No capability existed for conveniently submitting production to a central mainframe from distributed processors or intelligent terminals. A simple network status display was the only source of network status information available to operators; even less was available to users. In addition, network modeling efforts had to rely solely on individual machine performance measurements for network simulations and workload characterization. The FOCUS project offers a centralized solution to these problems.

IV. CONCEPTUAL REQUIREMENTS

A FOCUS analysis team was formed with representatives from operating systems development, network engineering, operations, and performance measurement staffs. The analysis team studied the current mainframe production subsystems and the ICN development plans. Production users and operations personnel were surveyed to determine current limitations and desirable new features to be included in the new design. These surveys were somewhat like "brainstorming" sessions--all ideas were heard <u>unimpeded</u> by feasibility or schedule considerations. This resulted in a plethora of suggestions from which a cohesive subset was selected to be implemented.

Implementation is being done in a phased manner: Phase 1 included a subset of functions to demonstrate the feasibility and usefulness of the project; Phase 2 includes full-featured enhancements to the Phase 1 system. For example, production control was implemented only for the Cray-1 supercomputers in Phase 1, with production control for CDC 7600s deferred to Phase 2. Requirements were identified only for the first two phases, leaving the third and later phases to be defined later.

V. DETAILED ANALYSIS

With conceptual requirements set, the analysis team explored structured methodologies for detailed analysis and design. The same input-transform-output structures that lend Jackson design techniques[5] to business applications seemed hard to apply to the interactive requirements of FOCUS. We felt a more complete and formal methodology than that offered by Dijkstra[6] was needed. We were also less familiar with the approaches of both Jackson and Dijkstra.

The structured approach based on the structured methodology of Yourdon and DeMarco was eventually chosen for several reasons. Foremost, its use had recently resulted in a very successful implementation of the Common File System Moreover, its emphasis on functional modularity meshed with our desire to use existing protocol modules and to port copies of the slave batch subsystems to mainframes and distributed processors as well as to intelligent terminals.

Delivering FOCUS on schedule was important, but the highest priority was reliability. Any inability to schedule work on four Cray-1 supercomputers for even a few hours would be prohibitively expensive. The "inch pebbles" instead of "milestones" approach offered an opportunity for better monitoring of the project, and use of proven protocol modules reduced the amount of work to be done, while it aided portability.

Figure 2 is an abstraction of the FOCUS system based on the top level data flow diagram of the new logical model. Performance measurement, production control, and information subsystems are all represented in their functional parts. Each subsystem consists of a central manager and various supporting subfunctions.

Use of such a logical model proved valuable because it allowed the analysis to concentrate on functionality, postponing implementation decisions until after the functional details were determined. That flexibility later resulted in an implementation of supporting utilities that substantially improved their portability to new operating systems.

Developing a structured specification from data-flow diagrams was also quite beneficial. This project involved many organizations, and the diagrams provided a common graphical perception of what was being developed. The analysis team was surprised at the communication issues and functional details resolved while developing the data-flow diagrams, data dictionary, and transform descriptions for the structured specification. Abstractions of these diagrams have also been used effectively for project reviews and operator learning.

VI. THREE SUBSYSTEMS

From the structured specification, the three main FOCUS subsystems were developed: performance measurement, production control, and network information.

The function of the performance measurement subsystem is to collect status and activity data from network nodes. These data are made available to the other two subsystems as a network status table. The information is also reduced and archived using a database management system for use in later network simulation and workload characterization efforts.

The production control subsystem manages production jobs for mainframe processors in the ICN. A production utility exists on each processor in the network that allows users on that processor to submit production jobs to FOCUS.

The centralized FOCUS production controller schedules all jobs by instructing the slave batch module on an appropriate mainframe processor to obtain and initiate the job.

The information subsystem is the window into the FOCUS system. Levels of information are available based on authorizations in a user permit table. Users can obtain status and activity of all network nodes and information about their own production jobs. Operators have further access to all production jobs, as well as online procedures and documentation. Information is transmitted as text displays except for color terminals where users have an option to receive data in a color graphics format.

VII. MODULARITY

A complex system is almost impossible to develop or maintain without dividing it into components of a manageable size. This may be the most important concept in software engineering; it has certainly been the cornerstone of the development of the Los Alamos ICN. It is the reason for our development of specialized satellite functions such as the Common File System and the Print and Graphics Express Station, as well as for FOCUS itself.

With such a network of functional nodes, communication between processes on different nodes is a key element. Because the network is comprised of equipment from many vendors, a set of hierarchical protocols has been developed locally for message and file transport and intermachine process-to-process communications. As equipment from a new vendor is integrated into the network, link drivers are written on a machine-dependent basis for communications. A higher-level communications package consisting of simple intermachine protocol (SIMP) and process-to-process (PTP) modules are installed to provide standardized communication throughout the network. Any capability developed using a processor for which SIMP and PTP already exist requires the development of the application software only, with the potentially difficult communications software available and, we hope, debugged. Figure 3 illustrates the communications protocol levels now in use at Los Alamos.

The FOCUS user interface, the utilities that reside on every processor and communicate with the FOCUS node, also benefit from design modularity. The importance of that interface to users was realized early, and care was taken to include user input in the design of that interface. As the system matures, however, that interface will evolve as new features are added and as experience using the system suggests improvements. These multiple-implementation utilities differ in order to match them to the processors on which they run. Implementations of such utilities for existing systems contained all parsing and decoding in each copy of the utility. The FOCUS design leaves only a processor-dependent shell in each copy that sets up communication and passes text in a standard format to a central utility implemented on the FOCUS node itself. Users' communication paths are actually lengthened but they gain substantial benefit from the common interface and higher reliability of a single, central module.

VIII. CONCLUSIONS

Software engineering practices and thorough analysis can result in timely development of large reliable computer systems like the Los Alamos Integrated Computing Network. With a foundation of functional communication modules, a network composed of powerful mainframe processors and specialized function nodes can be assembled The resulting modularity eases the difficult problem of integrating new hardware and capabilities into the network and adds strength to the implementation.

REFERENCES

1. Richard J. Wolfe, "A Satellite Output Graphics Station," Los Alamos National Laboratory report LA-UR-80-3716 (December 1980).

2. William W. Collins, Marjorie J. Devaney, and Emily W. Willbanks, "A Network File Storage System," Los Alamos National Laboratory report LA-UR-81-3558 (December 1981).

3. E. Yourdon and L. L. Constantine, Structured Design: Fundamentals of a Discipline of Computer Program and Systems Design (Prentice-Hall, Englewood Cliffs, N.J., 1979).

4. T. DeMarco, Structured Analysis and System Specification (Yourdon Press, New York, 1978).

5. M. Jackson, Principles of Program Design (Academic Press, New York, 1975).

6. J. Dijkstra, A Discipline of Programming (Prentice-Hall, Englewood Cliffs, N.J., 1976).

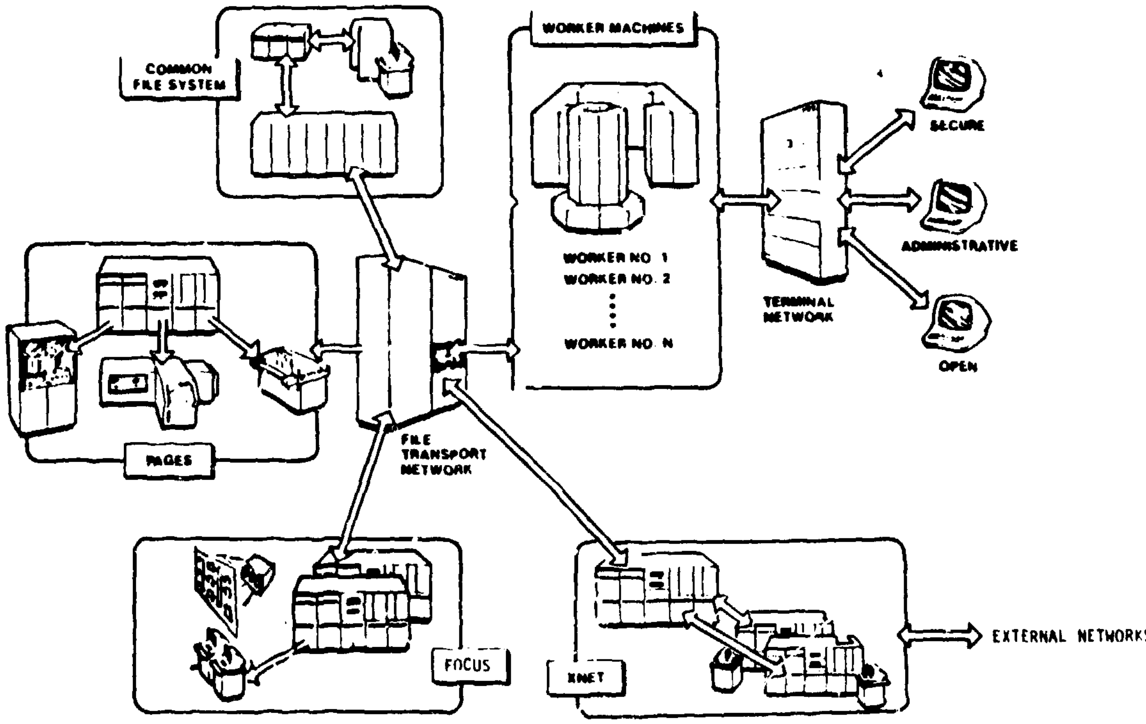

Fig. 1. Los Alamos National Laboratory Integrated Computing Network.

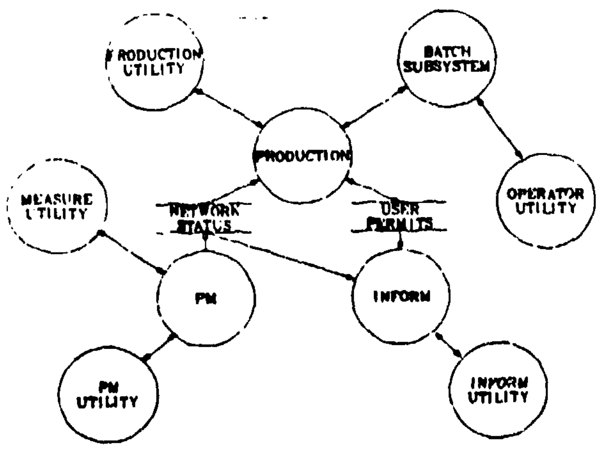

Fig. 2. FOCUS system.

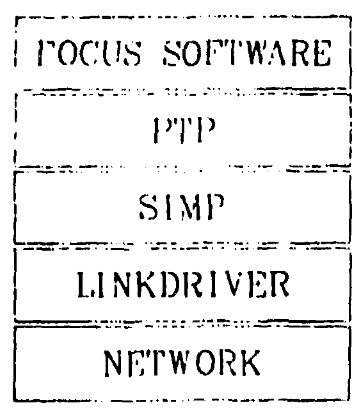

Fig. 3. Communications protocol levels.